

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 22-06-2018	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 8-Aug-2014 - 6-Feb-2018
---	--------------------------------	---

4. TITLE AND SUBTITLE Final Report: 1.3.2: Ergodic Control for Optimal Information Acquisition	5a. CONTRACT NUMBER W911NF-14-1-0461
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611102

6. AUTHORS	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Northwestern University Evanston Campus 1801 Maple Avenue Evanston, IL 60201 -3149	8. PERFORMING ORGANIZATION REPORT NUMBER
--	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 65339-NS.12

12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.
--

13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:	17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Todd Murphey
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU	19b. TELEPHONE NUMBER 847-467-1041

RPPR Final Report

as of 25-Jun-2018

Agency Code:

Proposal Number: 65339NS

Agreement Number: W911NF-14-1-0461

INVESTIGATOR(S):

Name: Todd David Murphey
Email: t-murphey@northwestern.edu
Phone Number: 8474671041
Principal: Y

Organization: **Northwestern University Evanston Campus**

Address: 1801 Maple Avenue, Evanston, IL 602013149

Country: USA

DUNS Number: 160079455

EIN: 362167817

Report Date: 06-May-2018

Date Received: 22-Jun-2018

Final Report for Period Beginning 08-Aug-2014 and Ending 06-Feb-2018

Title: 1.3.2: Ergodic Control for Optimal Information Acquisition

Begin Performance Period: 08-Aug-2014

End Performance Period: 06-Feb-2018

Report Term: 0-Other

Submitted By: Todd Murphey

Email: t-murphey@northwestern.edu

Phone: (847) 467-1041

Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 0

STEM Participants: 1

Major Goals: The objective of the proposed research is to develop an active search process that incorporates information measures into trajectory design to estimate the state of continuous, and potentially time-varying fields. Nonlinear optimal control theory plays a central role, and we are developing the tools necessary to select continuous-time trajectories for systems with nontrivial dynamics—for instance, systems with nonholonomic constraints, significant inertial effects, nonsmooth behavior—and often in high dimensions. The developed techniques use measures of ergodicity and the associated ergodic control to robustly explore for information while ensuring coverage of the informative subset of the state space (as indicated by prior measurements or even the lack thereof). The developed methods provide a synthesis technique for active sensing so that autonomous systems can optimize knowledge of continuous phenomena such as fluid or air flow fields, chemical dispersion, magnetic fields, electric fields, and other continuum phenomena. Further, by developing trajectories that enable fast and accurate estimation of the state and evolution of such fields in local regions, the research can support forecasting and prediction. Lastly, the outcomes of this work were tested both in computation and in experimental testbeds.

Accomplishments: The major accomplishments of this research can be summed up as the following.

1. The ability to compute near-optimal ergodic control for high dimensional, nonlinear systems in real-time, including stability guarantees expressed directly in terms of the distributions and experimental validation;
2. the ability to use ergodic control and information maximization to actively acquire data for Koopman operators—a spectral representation of a model that predicts a given set of data— including substantial real-time experimental validation.
3. alternative formulations of ergodic control that do not rely on spectral representations and instead use sample-based representations for control.

These outcomes are detailed below.

Ergodic Control

The research started with basic explorations into what ergodic control can accomplish. We showed that ergodic

RPPR Final Report as of 25-Jun-2018

exploration with known measurement models successfully rejects distractions and provides global guarantees regarding discovery in a given region [6]. We demonstrated that the ergodic metric can be used for a wide variety of nonlinear sensors on automation with nonlinear dynamics [7]. Moreover, we demonstrated that these techniques can be used in hardware, illustrating the method using an underwater robot that uses electrosense to localize conductors. (This application of near-field sensing enabled by ergodic control is relevant to many Army needs, including detection and localization of improvised explosive devices using near-field electromagnetic sensors.) Lastly, we showed how the trajectory optimization techniques can be extended to stochastic dynamics [3]. Although the exploration capabilities afforded by ergodic control appear powerful in terms of robust active sensing, the computations are substantial. We extended the ergodic control results to the context of real-time ergodic control [5], where we use a form of model-predictive control techniques to compute control updates quickly. We show there that the key to such an approach is satisfying a stability requirement with respect to both the reference distribution and the dynamical system's state. To generate a stability requirement that can be validated in real-time, we generalized sequential action control—a new method of approximate optimal control—so that it can use the ergodic metric as an objective function, leading to real-time control realizable for general nonlinear and hybrid systems. We used this control method to develop stability guarantees for ergodic control, where the stability guarantee uses the ergodic metric as a Lyapunov function on the distribution the trajectory generates, rather than the trajectory itself. As a result, we can provide analytical guarantees on long-time convergence of the controlled trajectory to the desired distribution (if the distribution is static). We have used this technique in a variety of simulated and experimental settings. These scenarios include simulated multi-agent tracking of an unknown number of targets with unknown dynamics, using a rotor vehicle dynamic model of each vehicle; this is a challenging and relevant problem to the army. Moreover, we applied the techniques experimentally to a single robot tracking multiple targets in our laboratory.

Spectral Representations of Dynamics

Exploration for localization and target identification/tracking is a useful application of ergodic control, but we also wanted to demonstrate that data-driven learning can also enable autonomy to adapt to its environment. Using the Koopman operator formalism, we showed that online learning of the nonparametric model of a robot is possible and substantially improves control performance [1]. In this paper, we developed a variation of model-based control for data-driven models, using sequential action control as the control methodology. The idea is very simple—we augment the state space with the control states, compute the Koopman operator that includes the model of the relationship between inputs and states, and then wrap model-based control around the result. Importantly, this all happens in real time. We used a simple robotic experiment of a robot on granular media to show that fast exploration yields an actionable control model that enables the robot to aggressively maneuver. Initially, with its nominal model, the robot was extremely ineffective at pursuing aggressive trajectories. After the Koopman model had been updated, the robot was able to execute much more aggressive trajectories, and did so while continuing to update its internal, data-driven model. In simulation, we showed that a rotor vehicle that has lost a rotor can arrest its fall in half the distance if it first takes one second to actively re-learn its own dynamics by maximizing information about the unknown model.

The Koopman formalism also provides a mechanism for active identification. While developing the results in 2), we quickly learned that the data obtained while controlling the robot in sand was insufficient to compute the Koopman operator; how one controls the robot has a dramatic impact on the quality of the spectrum obtained during operation. As a result, as a practical matter, a student then manually controlled the robot to get meaningful data. To automate the active construction of the operator, we initially wanted to use ergodic control to maximize the likelihood of obtaining good data, but we found that the calculations were prohibitive. However, a reasonably natural alternative turned out to be feasible. Maximizing the Fisher Information with respect to the infinite-dimensional spectrum is both computable as a trajectory optimization problem and has a natural interpretation in terms of the update to the operator itself [2]. Examples that we have completed include the robot operating in a granular environment, an aerial vehicle identifying its own dynamics due to an actuator failure, and a robot learning the dynamics of a pendulum inversion problem, for inversion and balance. All of these examples perform well only when active sensing, using the formal approaches developed as part of this work, are employed.

Future work should consider whether ergodic coverage can be used to generate reliable non-parametric models. Although we certainly see this in practice, we do not yet have a formal mathematical reason to believe that ergodic coverage should outperform information maximization with respect to the unknown operator.

Other Results may be found in our final written report.

[1] I. Abraham, G. de la Torre, and T. Murphey. Model-based control using Koopman operators. In *Robotics: Science and Systems Proceedings*, 2017.

[2] I. Abraham and T. D. Murphey. Active learning of dynamics and data-driven control. *IEEE Transactions on*

RPPR Final Report as of 25-Jun-2018

Robotics, Submitted.

- [3] G. De La Torre, K. Flaßkamp, A. Prabhakar, and T. D. Murphey. Ergodic exploration with stochastic sensor dynamics. In American Controls Conf. (ACC), pages 2971 – 2976, 2016.
- [4] G. De La Torre and T. D. Murphey. On the benefits of surrogate Lagrangians in optimal control and planning algorithms. In IEEE Int. Conf. on Decision and Control (CDC), pages 7384–7391, 2016.
- [5] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey. Real-time area coverage and target localization using receding-horizon ergodic exploration. IEEE Transactions on Robotics, 34(1):62–80, 2018.
- [6] L. Miller and T. D. Murphey. Optimal planning for target localization and coverage using range sensing. In IEEE Int. Conf. on Automation Science and Engineering (CASE), pages 501–508, 2015.
- [7] L. Miller, Y. Silverman, M. A. MacIver, and T. Murphey. Ergodic exploration of distributed information. IEEE Transactions on Robotics, 32(1):36–52, 2016.
- [8] A. Pervan and T. D. Murphey. Materials for Autonomous Systems, chapter Algorithmic Materials: Embedding computation within material properties as a means of enabling next generation autonomy. Elsevier, In Press. Eds. M. Strano and S. Walsh.

Training Opportunities: One graduate student--Taosha Fan (Ph.D., current, supported for one quarter)--were supported on this effort during the last months of the grant.

Results Dissemination: Eight publications were produced as part of this work.

- [1] I. Abraham, G. de la Torre, and T. Murphey. Model-based control using Koopman operators. In Robotics: Science and Systems Proceedings, 2017.
- [2] I. Abraham and T. D. Murphey. Active learning of dynamics and data-driven control. IEEE Transactions on Robotics, Submitted.
- [3] G. De La Torre, K. Flaßkamp, A. Prabhakar, and T. D. Murphey. Ergodic exploration with stochastic sensor dynamics. In American Controls Conf. (ACC), pages 2971 – 2976, 2016.
- [4] G. De La Torre and T. D. Murphey. On the benefits of surrogate Lagrangians in optimal control and planning algorithms. In IEEE Int. Conf. on Decision and Control (CDC), pages 7384–7391, 2016.
- [5] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey. Real-time area coverage and target localization using receding-horizon ergodic exploration. IEEE Transactions on Robotics, 34(1):62–80, 2018.
- [6] L. Miller and T. D. Murphey. Optimal planning for target localization and coverage using range sensing. In IEEE Int. Conf. on Automation Science and Engineering (CASE), pages 501–508, 2015.
- [7] L. Miller, Y. Silverman, M. A. MacIver, and T. Murphey. Ergodic exploration of distributed information. IEEE Transactions on Robotics, 32(1):36–52, 2016.
- [8] A. Pervan and T. D. Murphey. Materials for Autonomous Systems, chapter Algorithmic Materials: Embedding computation within material properties as a means of enabling next generation autonomy. Elsevier, In Press. Eds. M. Strano and S. Walsh.

Honors and Awards: The PI was a member of the National Academies / National Research Council Committee on Counter- Unmanned Aircraft System (CUAS) Capability for Battalion-and-Below Operations (2016).

Protocol Activity Status:

Technology Transfer: Nothing to Report

PARTICIPANTS:

Participant Type: PD/PI

Participant: Todd Murphey

Person Months Worked: 1.00

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Funding Support:

Participant Type: Graduate Student (research assistant)

RPPR Final Report

as of 25-Jun-2018

Participant: Taosha Fan

Person Months Worked: 5.00

Project Contribution:

International Collaboration:

International Travel:

National Academy Member: N

Other Collaborators:

Funding Support:

ARTICLES:

Publication Type: Journal Article

Peer Reviewed: Y

Publication Status: 1-Published

Journal: IEEE Transactions on Robotics

Publication Identifier Type: DOI

Publication Identifier: 10.1109/TRO.2017.2766265

Volume: 34

Issue: 1

First Page #: 62

Date Submitted: 6/22/18 12:00AM

Date Published: 2/1/18 6:00AM

Publication Location:

Article Title: Real-Time Area Coverage and Target Localization Using Receding-Horizon Ergodic Exploration

Authors: Anastasia Mavrommati, Emmanouil Tzorakoleftherakis, Ian Abraham, Todd D. Murphey

Keywords: Mobile robots, motion planning, nonlinear control systems, unmanned autonomous vehicles

Abstract: Although a number of solutions exist for the problems of coverage, search, and target localization- commonly addressed separately-whether there exists a unified strategy that addresses these objectives in a coherent manner without being application specific remains a largely open research question. In this paper, we develop a receding-horizon ergodic control approach, based on hybrid systems theory, that has the potential to fill this gap. The nonlinear model-predictive control algorithm plans real-time motions that optimally improve ergodicity with respect to a distribution defined by the expected information density across the sensing domain. We establish a theoretical framework for global stability guarantees with respect to a distribution. Moreover, the approach is distributable across multiple agents so that each agent can independently compute its own control while sharing statistics of its coverage across a communication network. We demonstrate the method in both simulation and i

Distribution Statement: 1-Approved for public release; distribution is unlimited.

Acknowledged Federal Support: Y

CONFERENCE PAPERS:

Publication Type: Conference Paper or Presentation

Publication Status: 1-Published

Conference Name: IEEE Int. Conf. on Automation Science and Engineering

Date Received: 25-Aug-2016

Conference Date: 24-Aug-2015

Date Published:

Conference Location: Gothenberg, Sweden

Paper Title: Optimal Planning for Target Localization and Coverage Using Range Sensing

Authors: Lauren Miller, Todd Murphey

Acknowledged Federal Support: Y

Publication Type: Conference Paper or Presentation

Publication Status: 1-Published

Conference Name: 2016 American Control Conference (ACC)

Date Received: 25-Aug-2016

Conference Date: 06-Jul-2016

Date Published:

Conference Location: Boston, MA, USA

Paper Title: Ergodic exploration with stochastic sensor dynamics

Authors: Gerardo De La Torre, Kathrin Flaßkamp, Ahalya Prabhakar, Todd D. Murphey

Acknowledged Federal Support: Y

RPPR Final Report
as of 25-Jun-2018

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: IEEE Conference on Decision and Control
Date Received: 31-Aug-2017 Conference Date: 17-Dec-2016 Date Published: 25-Aug-2016
Conference Location: Las Vegas, USA
Paper Title: On the Benefits of Surrogate Lagrangians in Optimal Control and Planning Algorithms
Authors: Gerardo De La Torre, Todd D. Murphey
Acknowledged Federal Support: **Y**

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: Robotics: Science and Systems
Date Received: 22-Jun-2018 Conference Date: 15-Jul-2017 Date Published: 01-Jul-2017
Conference Location: Boston, MA, USA
Paper Title: Model-based control using Koopman operators
Authors: I. Abraham, G. de la Torre, and T. Murphey
Acknowledged Federal Support: **Y**

DISSERTATIONS:

Publication Type: Thesis or Dissertation
Institution: Northwestern University
Date Received: 25-Aug-2016 Completion Date: 12/1/15 6:04PM
Title: Optimal Ergodic Control for Active Search and Information Acquisition
Authors: Lauren Miller
Acknowledged Federal Support: **N**

Publication Type: Thesis or Dissertation
Institution: Northwestern University
Date Received: 31-Aug-2017 Completion Date: 6/1/17 5:21PM
Title: Real-Time Algorithms for Symbol-Based Automation
Authors: Anastasia Mavrommati
Acknowledged Federal Support: **Y**

Final Report for Grant W911NF1410461

1.3.2: Ergodic Control for Optimal Information Acquisition

Todd D. Murphey
Northwestern University

1 Major Goals

The objective of the proposed research is to develop an active search process that incorporates information measures into trajectory design to estimate the state of continuous, and potentially time-varying fields. Nonlinear optimal control theory plays a central role, and we are developing the tools necessary to select continuous-time trajectories for systems with nontrivial dynamics—for instance, systems with nonholonomic constraints, significant inertial effects, nonsmooth behavior—and often in high dimensions. The developed techniques use measures of ergodicity and the associated ergodic control to robustly explore for information while ensuring coverage of the informative subset of the state space (as indicated by prior measurements or even the lack thereof). The developed methods provide a synthesis technique for active sensing so that autonomous systems can optimize knowledge of continuous phenomena such as fluid or air flow fields, chemical dispersion, magnetic fields, electric fields, and other continuum phenomena. Further, by developing trajectories that enable fast and accurate estimation of the state and evolution of such fields in local regions, the research can support forecasting and prediction. Lastly, the outcomes of this work were tested both in computation and in experimental testbeds.

2 Accomplishments

The major accomplishments of this research can be summed up as the following.

1. The ability to compute near-optimal ergodic control for high dimensional, nonlinear systems in real-time, including stability guarantees expressed directly in terms of the distributions and experimental validation;
2. the ability to use ergodic control and information maximization to active acquire data for Koopman operators—a spectral representation of a model that predicts a given set of data—including substantial real-time experimental validation.
3. alternative formulations of ergodic control that do not rely on spectral representations and instead use sample-based representations for control.

These outcomes are detailed below.

2.1 Ergodic Control

The research started with basic explorations into what ergodic control can accomplish. We showed that ergodic exploration with known measurement models successfully rejects distractions and provides global guarantees regarding discovery in a given region [6]. We demonstrated that the ergodic

metric can be used for a wide variety of nonlinear sensors on automation with nonlinear dynamics [7]. Moreover, we demonstrated that these techniques can be used in hardware, illustrating the method using an underwater robot that uses electrosense to localize conductors. (This application of near-field sensing enabled by ergodic control is relevant to many Army needs, including detection and localization of improvised explosive devices using near-field electromagnetic sensors.) Lastly, we showed how the trajectory optimization techniques can be extended to stochastic dynamics [3].

Although the exploration capabilities afforded by ergodic control appear powerful in terms of robust active sensing, the computations are substantial. We extended the ergodic control results to the context of real-time ergodic control [5], where we use a form of model-predictive control techniques to compute control updates quickly. We show there that the key to such an approach is satisfying a stability requirement with respect to both the reference distribution and the dynamical system's state. To generate a stability requirement that can be validated in real-time, we generalized sequential action control—a new method of approximate optimal control—so that it can use the ergodic metric as an objective function, leading to real-time control realizable for general nonlinear and hybrid systems. We used this control method to develop stability guarantees for ergodic control, where the stability guarantee uses the ergodic metric as a Lyapunov function on the distribution the trajectory generates, rather than the trajectory itself. As a result, we can provide analytical guarantees on long-time convergence of the controlled trajectory to the desired distribution (if the distribution is static). We have used this technique in a variety of simulated and experimental settings. These scenarios include simulated multi-agent tracking of an unknown number of targets with unknown dynamics, using a rotor vehicle dynamic model of each vehicle; this is a challenging and relevant problem to the army. Moreover, we applied the techniques experimentally to a single robot tracking multiple targets in our laboratory.

Although the student was not supported financially by this grant, we were able to show that ergodic control is naturally distributable. That is, standard consensus algorithms are sufficient to distribute ergodic control over an arbitrary number of agents. Although not of particular theoretical interest, this is a result that likely has practical value.

2.2 Spectral Representations of Dynamics

Exploration for localization and target identification/tracking is a useful application of ergodic control, but we also wanted to demonstrate that data-driven learning can also enable autonomy to adapt to its environment. Using the Koopman operator formalism, we showed that online learning of the nonparametric model of a robot is possible and substantially improves control performance [1]. In this paper, we developed a variation of model-based control for data-driven models, using sequential action control as the control methodology. The idea is very simple—we augment the state space with the control states, compute the Koopman operator that includes the model of the relationship between inputs and states, and then wrap model-based control around the result. Importantly, this all happens in real time. We used a simple robotic experiment of a robot on granular media to show that fast exploration yields an actionable control model that enables the robot to aggressively maneuver. Initially, with its nominal model, the robot was extremely ineffective at pursuing aggressive trajectories. After the Koopman model had been updated, the robot was able to execute much more aggressive trajectories, and did so while continuing to update its internal, data-driven model. In simulation, we showed that a rotorvehicle that has lost a rotor can arrest

its fall in half the distance if it first takes one second to actively re-learn its own dynamics by maximizing information about the unknown model.

The Koopman formalism also provides a mechanism for active identification. While developing the results in 2), we quickly learned that the data obtained while controlling the robot in sand was insufficient to compute the Koopman operator; how one controls the robot has a dramatic impact on the quality of the spectrum obtained during operation. As a result, as a practical matter, a student then manually controlled the robot to get *meaningful* data. To automate the active construction of the operator, we initially wanted to use ergodic control to maximize the likelihood of obtaining good data, but we found that the calculations were prohibitive. However, a reasonably natural alternative turned out to be feasible. Maximizing the Fisher Information with respect to the infinite-dimensional spectrum is both computable as a trajectory optimization problem and has a natural interpretation in terms of the update to the operator itself [2]. Examples that we have completed include the robot operating in a granular environment, an aerial vehicle identifying its own dynamics due to an actuator failure, and a robot learning the dynamics of a pendulum inversion problem, for inversion and balance. All of these examples perform well only when active sensing, using the formal approaches developed as part of this work, are employed.

Future work should consider whether ergodic coverage can be used to generate reliable non-parametric models. Although we certainly see this in practice, we do not yet have a formal mathematical reason to believe that ergodic coverage should outperform information maximization with respect to the unknown operator.

2.3 Other Results

Two other subjects were actively pursued near the end of this work: alternative metrics for ergodic control and the connection between information-based control and “algorithmic” materials.

Ergodic Control, Occupation Measures, and Kullback-Leibler Divergence: One of the problems with ergodic control, particularly with respect to computation, is that it relies on the Fourier decomposition of both a reference density function as well as the trajectory itself. This is very computationally intensive, particularly as the dimension of space over which the density function takes values grows. In robotics, this is exactly where sample-based methods have come to dominate practical implementation.

Initially, we looked at the occupation-measure formulation of ergodic control, since it has been used in other areas of spectral approaches to control. We were able to compute ergodic trajectories this way, but the computations were not particularly faster and were substantially more brittle than those that use the Fourier decomposition.

However, an idea close in spirit to occupation measures was to treat ergodic control as a requirement on the Kullback-Leibler divergence between two distributions. Ergodicity can be interpreted as the Kullback-Leibler distance as evaluated for a trajectory, rather than a distribution that has full support over the same space as the reference distribution; the Kullback-Leibler divergence will always be infinite between a distribution and a trajectory. If we replace a point in the trajectory with a volume—like a time-varying occupation measure—the K-L divergence is finite and can be approximated using sample-based methods. We find that this is easily executable so long as the ergodic trajectory does not need to reproduce fine-scale detail in the original reference distribution.

(This work has not yet been published, but the PI is supporting it on another grant where the results can be used.)

Algorithmic Materials: Lastly, the end of this grant included investigation into the development of algorithmic materials. Our initial work can be found [8], where we include preliminary results on the relationship between control and estimation theory and the physical algorithmic capabilities that materials may provide.

First, it is worth noting that at very small scales, traditional computation will often be impossible. Nevertheless, sensing, actuation, memory, and operations on these signals may be available. The question we were interested in is whether one can translate control techniques to the design of these materials and their interconnections. To do so, we looked at a process that starts with a task description in the form of an objective function, approximates an in-principle control policy that accomplishes the task, and then projects that policy onto feasible physical designs. We showed in an example of a Micro State Machine—a model motivated by a graphene body with simple chemical components such as chemical comparators for sensing and potential poisoning for actuation—that a finite-state machine that involves only feasible components leads to a physical design that can move the body to a desired location. Moreover, we also showed that some tasks *cannot* be achieved this way; specifically, ergodic coverage cannot be achieved unless one has access to some form of memory. Regardless, this work represents our first steps towards a theory of *cyber-free* autonomy, specifically at microscales.

References

- [1] I. Abraham, G. de la Torre, and T. Murphey. Model-based control using Koopman operators. In *Robotics: Science and Systems Proceedings*, 2017.
- [2] I. Abraham and T. D. Murphey. Active learning of dynamics and data-driven control. *IEEE Transactions on Robotics*, Submitted.
- [3] G. De La Torre, K. Flaßkamp, A. Prabhakar, and T. D. Murphey. Ergodic exploration with stochastic sensor dynamics. In *American Controls Conf. (ACC)*, pages 2971 – 2976, 2016.
- [4] G. De La Torre and T. D. Murphey. On the benefits of surrogate Lagrangians in optimal control and planning algorithms. In *IEEE Int. Conf. on Decision and Control (CDC)*, pages 7384–7391, 2016.
- [5] A. Mavrommati, E. Tzorakoleftherakis, I. Abraham, and T. D. Murphey. Real-time area coverage and target localization using receding-horizon ergodic exploration. *IEEE Transactions on Robotics*, 34(1):62–80, 2018.
- [6] L. Miller and T. D. Murphey. Optimal planning for target localization and coverage using range sensing. In *IEEE Int. Conf. on Automation Science and Engineering (CASE)*, pages 501–508, 2015.
- [7] L. Miller, Y. Silverman, M. A. MacIver, and T. Murphey. Ergodic exploration of distributed information. *IEEE Transactions on Robotics*, 32(1):36–52, 2016.
- [8] A. Pervan and T. D. Murphey. *Materials for Autonomous Systems*, chapter Algorithmic Materials: Embedding computation within material properties as a means of enabling next generation autonomy. Elsevier, In Press. Eds. M. Strano and S. Walsh.